L Number	Hits	Search Text	DB	Time stamp
1	48247	out adj2 order	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 10:24
2	86416	cach$3	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 10:25
3	1161	(bank or module) near5 conflict$3	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 10:44
4	12969	(bank or module) near5 (overlap$4 or over-lap$4 or simultaneous$2)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 10:27
5	54	(out adj2 order) same ((bank or module) near5 conflict$3)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 10:27
7	3695	(queue or fifo) near3 (pend$3 or await$3 or outstand$3)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 10:45
8	4	((bank or module) near5 (overlap$4 or over-lap$4 or simultaneous$2)) same (out adj2 order)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 10:29
9	2	cach$3 same (((bank or module) near5 (overlap$4 or over-lap$4 or simultaneous$2)) same (out adj2 order))	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 10:29
10	10	((queue or fifo) near3 (pend$3 or await$3 or outstand$3)) and (cach$3 same ((out adj2 order) same ((bank or module) near5 conflict$3)))	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 10:31
11	13937	no$1 adj block$3	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 10:32
12	398	(no$1 adj block$3) with cach$3	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 10:44
13	48	((queue or fifo) near3 (pend$3 or await$3 or outstand$3)) and ((no$1 adj block$3) with cach$3)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 10:32
14	3981	(bank or module) near5 cache	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 10:33
15	18	((bank or module) near5 cache) and (((queue or fifo) near3 (pend$3 or await$3 or outstand$3)) and ((no$1 adj block$3) with cach$3))	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 10:35

16	60	((bank or module) near5 conflict$3) and ((queue or fifo) near3 (pend$3 or await$3 or outstand$3))	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 10:37
17	211	cach$3 same ((bank or module) near5 conflict$3)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 10:37
18	28	(cach$3 same ((bank or module) near5 conflict$3)) and (((bank or module) near5 conflict$3) and ((queue or fifo) near3 (pend$3 or await$3 or outstand$3)))	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 10:40
19	110	((no$1 adj block$3) with cach$3) and ((bank or module) near5 cache)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 10:40
20	18	((queue or fifo) near3 (pend$3 or await$3 or outstand$3)) and (((no$1 adj block$3) with cach$3) and ((bank or module) near5 cache))	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 10:40
6	51	cach$3 same ((out adj2 order) same ((bank or module) near5 conflict$3))	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 10:41
21	110	((bank or module) near5 cache) and ((no$1 adj block$3) with cach$3)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 10:45
22	90	(((no$1 adj block$3) with cach$3) and ((bank or module) near5 cache)) not (cach$3 same ((out adj2 order) same ((bank or module) near5 conflict$3)))	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 10:45
23	43	(queue or fifo) and ((((no$1 adj block$3) with cach$3) and ((bank or module) near5 cache)) not (cach$3 same ((out adj2 order) same ((bank or module) near5 conflict$3))))	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 10:45

L Number	Hits	Search Text	DB	Time stamp
1	2	6711654.pn.	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 13:42
2	4	("5740402" | "6081873" | "6425044" | "6539457").PN.	USPAT	2004/06/28 13:34
3	0	conflict* with bank	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 13:42
4	554	conflict$3 near5 bank	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 13:42
5	106	(queue or fifo) same (conflict$3 near5 bank)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 13:43
6	2266	cache with bank	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 13:43
7	60	((queue or fifo) same (conflict$3 near5 bank)) and (cache with bank)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 13:49
8	104	(conflict$3 near5 bank) same pipelin$3	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 13:49
10	74	(load$3 or read$3 or fetch$3) same (stor$3 or writ$3) same ((conflict$3 near5 bank) same pipelin$3)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 13:52
11	68	cache and ((load$3 or read$3 or fetch$3) same (stor$3 or writ$3) same ((conflict$3 near5 bank) same pipelin$3))	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 13:53
12	50	(cache with bank) and ((load$3 or read$3 or fetch$3) same (stor$3 or writ$3) same ((conflict$3 near5 bank) same pipelin$3))	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 13:54
13	41022	(multiport or (multi or dual or two)) adj port$3	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 13:55
14	854	((multiport or (multi or dual or two)) adj port$3) near4 cach$3	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 13:55
15	216	(((multiport or (multi or dual or two)) adj port$3) near4 cach$3) and (cache with bank)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 13:55

16	92	(conflict$3 near5 bank) and ((((multiport or (multi or dual or two)) adj port$3) near4 cach$3) and (cache with bank))	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 13:56
17	38	((queue or fifo) same (conflict$3 near5 bank)) and ((((multiport or (multi or dual or two)) adj port$3) near4 cach$3) and (cache with bank))	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 13:56
18	48247	out adj2 order	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 13:58
19	54	(out adj2 order) same (conflict$3 near5 bank)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 13:58
20	40	((queue or fifo) same (conflict$3 near5 bank)) same (out adj2 order)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 13:59
21	1	port and 6711654.pn.	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/06/28 13:59

Welcome to IEEE Xplore®

- Home
- What Can I Access?
- Log-out

Tables of Contents

- Journals & Magazines
- Conference Proceedings
- Standards

Search

- By Author
- Basic
- Advanced

Member Services

- Join IEEE
- Establish IEEE Web Account
- Access the IEEE Member Digital Library

IEEE Enterprise

- Access the IEEE Enterprise File Cabinet

Print Format

Your search matched **20** of **1051129** documents.
A maximum of **500** results are displayed, **15** to a page, sorted by **Relevance** in **Descending** order.

Refine This Search:
You may refine your search by editing the current search expression or entering a new one in the text box.

bank <near/1> conflict* Search

☐ Check to search within this result set

Results Key:
JNL = Journal or Magazine **CNF** = Conference **STD** = Standard

---

1 **Architecture scalability of parallel vector computers with a shared memory**
Dekker, E.;
Computers, IEEE Transactions on , Volume: 47 , Issue: 5 , May 1998
Pages:614 - 624

[Abstract] [PDF Full-Text (168 KB)] **IEEE JNL**

---

2 **Multiskewing-a novel technique for optimal parallel memory access**
Deb, A.;
Parallel and Distributed Systems, IEEE Transactions on , Volume: 7 , Issue: 6 , June 1996
Pages:595 - 604

[Abstract] [PDF Full-Text (852 KB)] **IEEE JNL**

---

3 **A memory interference model for regularly patterned multiple stream vector accesses**
Qing Yang; Tao Yang;
Parallel and Distributed Systems, IEEE Transactions on , Volume: 6 , Issue: 5 , May 1995
Pages:520 - 530

[Abstract] [PDF Full-Text (968 KB)] **IEEE JNL**

---

4 **On memory contention problems in vector multiprocessors**
Fricker, C.;
Computers, IEEE Transactions on , Volume: 44 , Issue: 1 , Jan. 1995
Pages:92 - 105

[Abstract] [PDF Full-Text (1056 KB)] **IEEE JNL**

---

5 **Conflict free memory addressing for dedicated FFT hardware**
Johnson, L.G.;
Circuits and Systems II: Analog and Digital Signal Processing, IEEE Transactions on [see also Circuits and Systems II: Express Briefs, IEEE Transactions on] , Volume: 39 , Issue: 5 , May 1992

Pages:312 - 316

[Abstract] [PDF Full-Text (412 KB)] **IEEE JNL**

---

6 **A loop partition technique for reducing cache bank conflict in multithreaded architecture**
Wu, C.-C.; Chen, C.F.;
Computers and Digital Techniques, IEE Proceedings- , Volume: 143 , Issue: 1 , Jan. 1996
Pages:30 - 36

[Abstract] [PDF Full-Text (718 KB)] **IEE JNL**

---

7 **Design and implementation of high-performance memory systems for future packet buffers**
Garcia, J.; Corbal, J.; Cerda, L.; Valero, M.;
Microarchitecture, 2003. MICRO-36. Proceedings. 36th Annual IEEE/ACM International Symposium on , 3-5 Dec. 2003
Pages:372 - 384

[Abstract] [PDF Full-Text (473 KB)] **IEEE CNF**

---

8 **Resolving register bank conflicts for a network processor**
Zhuang, X.; Santosh Pande;
Parallel Architectures and Compilation Techniques, 2003. PACT 2003. Proceedings. 12th International Conference on , 27 Sept.-1 Oct. 2003
Pages:269 - 278

[Abstract] [PDF Full-Text (284 KB)] **IEEE CNF**

---

9 **System-level performance optimization of the data queueing memory management in high-speed network processors**
Ykman-Couvreur, C.; Lambrecht, J.; Verkest, D.; Catthoor, F.; Nikologiannis, A.; Konstantoulakis, G.;
Design Automation Conference, 2002. Proceedings. 39th , 10-14 June 2002
Pages:518 - 523

[Abstract] [PDF Full-Text (800 KB)] **IEEE CNF**

---

10 **Reducing register ports for higher speed and lower energy**
Park, I.; Powell, M.D.; Vijaykumar, T.N.;
Microarchitecture, 2002. (MICRO-35). Proceedings. 35th Annual IEEE/ACM International Symposium on , 18-22 Nov. 2002
Pages:171 - 182

[Abstract] [PDF Full-Text (277 KB)] **IEEE CNF**

---

11 **Efficient per-flow queueing in DRAM at OC-192 line rate using out-of-order execution techniques**
Nikologiannis, A.; Katevenis, M.;
Communications, 2001. ICC 2001. IEEE International Conference on , Volume: 7 , 11-14 June 2001
Pages:2048 - 2052 vol.7

[Abstract] [PDF Full-Text (540 KB)] **IEEE CNF**

---

12 **Parallel cachelets**
Limaye, D.; Rakvic, R.; Shen, J.P.;
Computer Design, 2001. ICCD 2001. Proceedings. 2001 International Conference on , 23-26 Sept. 2001

Pages:284 - 292

[Abstract] [PDF Full-Text (952 KB)] **IEEE CNF**

---

13 **A DRAM system for consistently reducing CPU wait cycles**
Kanno, Y.; Mizuno, H.; Watanabe, T.;
VLSI Circuits, 1999. Digest of Technical Papers. 1999 Symposium on , 17-19 June 1999
Pages:131 - 132

[Abstract] [PDF Full-Text (188 KB)] **IEEE CNF**

---

14 **Speculation techniques for improving load related instruction scheduling**
Yoaz, A.; Erez, M.; Ronen, R.; Jourdan, S.;
Computer Architecture, 1999. Proceedings of the 26th International Symposium on , 2-4 May 1999
Pages:42 - 53

[Abstract] [PDF Full-Text (212 KB)] **IEEE CNF**

---

15 **A memory system supporting the efficient SIMD computation of the two dimensional DWT**
Trenas, M.A.; Lopez, J.; Zapata, E.L.; Arguello, F.;
Acoustics, Speech, and Signal Processing, 1998. ICASSP '98. Proceedings of the 1998 IEEE International Conference on , Volume: 3 , 12-15 May 1998
Pages:1521 - 1524 vol.3

[Abstract] [PDF Full-Text (340 KB)] **IEEE CNF**

Your search matched **20** of **1051129** documents.
A maximum of **500** results are displayed, **15** to a page, sorted by **Relevance** in **Descending** order.

Refine This Search:
You may refine your search by editing the current search expression or entering a new one in the text box.

bank <near/1> conflict* [Search]

☐ Check to search within this result set

Results Key:
JNL = Journal or Magazine **CNF** = Conference **STD** = Standard

---

16 **Enhancing vector access performance in CRAY X-MP memory system**
Edirisooriya, S.; Edirisooriya, G.;
Compcon Spring '93, Digest of Papers. , 22-26 Feb. 1993
Pages:569 - 576

[Abstract] [PDF Full-Text (460 KB)] **IEEE CNF**

---

17 **A vector memory system based on wafer-scale integrated memory arrays**
Chiueh, T.-C.;
Computer Design: VLSI in Computers and Processors, 1993. ICCD '93. Proceedings., 1993 IEEE International Conference on , 3-6 Oct. 1993
Pages:284 - 288

[Abstract] [PDF Full-Text (456 KB)] **IEEE CNF**

---

18 **Memory bank conflict on some vector supercomputers**
Fujino, S.;
CompEuro '92 . 'Computer Systems and Software Engineering',Proceedings. , 4-8 May 1992
Pages:359 - 364

[Abstract] [PDF Full-Text (396 KB)] **IEEE CNF**

---

19 **A fine-grained MIMD architecture based upon register channels**
Gupta, R.;
Microprogramming and Microarchitecture. Micro 23. Proceedings of the 23rd Annual Workshop and Symposium., Workshop on , 27-29 Nov. 1990
Pages:28 - 37

[Abstract] [PDF Full-Text (736 KB)] **IEEE CNF**

---

20 **The impact of memory organization on the performance of matrix multiplication**
Hake, J.-F.; Homberg, W.;
Supercomputing '90. Proceedings of , 12-16 Nov. 1990
Pages:34 - 40

Your search matched **0** of **1047691** documents.
A maximum of **500** results are displayed, **15** to a page, sorted by **Relevance** in **Descending** order.

Refine This Search:
You may refine your search by editing the current search expression or entering a new one in the text box.

(no* <near/1> block*) <near/1> cach* Search

☐ Check to search within this result set

Results Key:
JNL = Journal or Magazine **CNF** = Conference **STD** = Standard

= Your access to full-text

Results:
No documents matched your query.

Try our New Full-text Search Prototype GO

Help

1) Enter a single keyword, phrase, or Boolean expression.
Example: acoustic imaging (means the phrase acoustic imaging plus any stem variations)
2) Limit your search by using search operators and field codes, if desired.
Example: optical <and> (fiber <or> fibre) <in> ti
3) Limit the results by selecting Search Options.
4) Click Search. See Search Examples

```
cach* and bank and conflict*
```

Start Search | Clear

Note: This function returns plural and suffixed forms of the keyword(s).

Search operators: <and> <or> <not> <in> More

Field codes: au (author), ti (title), ab (abstract), jn (publication name), de (index term) More

Search Options:

Select publication types:

- [x] IEEE Journals
- [x] IEE Journals
- [x] IEEE Conference proceedings
- [x] IEE Conference proceedings
- [x] IEEE Standards

Select years to search:

From year: All to Present

Organize search results by:

Sort by: Relevance

In: Descending order

List 15 Results per page

Your search matched **12** of **1047691** documents.
A maximum of **500** results are displayed, **15** to a page, sorted by **Relevance** in **Descending** order.

Refine This Search:
You may refine your search by editing the current search expression or entering a new one in the text box.

cach* and bank and conflict* Search

Check to search within this result set

Results Key:
JNL = Journal or Magazine **CNF** = Conference **STD** = Standard
= Your access to full-text

---

1 **An efficient instruction cache scheme for object-oriented languages**
Yul Chu; Ito, M.R.;
Performance, Computing, and Communications, 2001. IEEE International Conference on. , 4-6 April 2001
Pages:329 - 336

[Abstract] [PDF Full-Text (592 KB)] **IEEE CNF**

---

2 **A loop partition technique for reducing cache bank conflict in multithreaded architecture**
Wu, C.-C.; Chen, C.F.;
Computers and Digital Techniques, IEE Proceedings- , Volume: 143 , Issue: 1 , Jan. 1996
Pages:30 - 36

[Abstract] [PDF Full-Text (718 KB)] **IEE JNL**

---

3 **The 2-way thrashing-avoidance cache (TAC): an efficient instruction cache scheme for object-oriented languages**
Chu, Y.; Ito, M.R.;
Computer Design, 2000. Proceedings. 2000 International Conference on , 17-20 Sept. 2000
Pages:93 - 98

[Abstract] [PDF Full-Text (456 KB)] **IEEE CNF**

---

4 **Parallel cachelets**
Limaye, D.; Rakvic, R.; Shen, J.P.;
Computer Design, 2001. ICCD 2001. Proceedings. 2001 International Conference on , 23-26 Sept. 2001
Pages:284 - 292

[Abstract] [PDF Full-Text (952 KB)] **IEEE CNF**

---

5 **Speculation techniques for improving load related instruction scheduling**
Yoaz, A.; Erez, M.; Ronen, R.; Jourdan, S.;

Computer Architecture, 1999. Proceedings of the 26th International Symposium on , 2-4 May 1999
Pages:42 - 53

[Abstract] [PDF Full-Text (212 KB)] **IEEE CNF**

---

6 **A technique for high bandwidth and deterministic low latency load/store accesses to multiple cache banks**
Neefs, H.; Vandierendonck, H.; De Bosschere, K.;
High-Performance Computer Architecture, 2000. HPCA-6. Proceedings. Sixth International Symposium on , 8-12 Jan. 2000
Pages:313 - 324

[Abstract] [PDF Full-Text (284 KB)] **IEEE CNF**

---

7 **A vector memory system based on wafer-scale integrated memory arrays**
Chiueh, T.-C.;
Computer Design: VLSI in Computers and Processors, 1993. ICCD '93. Proceedings., 1993 IEEE International Conference on , 3-6 Oct. 1993
Pages:284 - 288

[Abstract] [PDF Full-Text (456 KB)] **IEEE CNF**

---

8 **A permutation-based page interleaving scheme to reduce row-buffer conflicts and exploit data locality**
Zhao Zhang; Zhichun Zhu; Xiaodong Zhang;
Microarchitecture, 2000. MICRO-33. Proceedings. 33rd Annual IEEE/ACM International Symposium on , 10-13 Dec. 2000
Pages:32 - 41

[Abstract] [PDF Full-Text (876 KB)] **IEEE CNF**

---

9 **Co-design of interleaved memory systems**
Hua Lin; Wolf, W.;
Hardware/Software Codesign, 2000. CODES 2000. Proceedings of the Eighth International Workshop on , 3-5 May 2000
Pages:46 - 50

[Abstract] [PDF Full-Text (480 KB)] **IEEE CNF**

---

10 **A comparative analysis of cache designs for vector processing**
Tong Sun; Qing Yang;
Computers, IEEE Transactions on , Volume: 48 , Issue: 3 , March 1999
Pages:331 - 344

[Abstract] [PDF Full-Text (360 KB)] **IEEE JNL**

---

11 **An object-oriented data cache architecture for programmable parallel digital signal processors**
Kneip, J.;
Algorithms and Architectures for Parallel Processing, 1997. ICAPP 97. 1997 3rd International Conference on , 10-12 Dec. 1997
Pages:105 - 112

[Abstract] [PDF Full-Text (532 KB)] **IEEE CNF**

---

12 **A DRAM system for consistently reducing CPU wait cycles**
Kanno, Y.; Mizuno, H.; Watanabe, T.;
VLSI Circuits, 1999. Digest of Technical Papers. 1999 Symposium on , 17-19 June

1999
Pages:131 - 132

[Abstract] [PDF Full-Text (188 KB)] **IEEE CNF**

PORTAL

US Patent & Trademark Office

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: ◉ The ACM Digital Library ○ The Guide

(bank <sentence> conflict*) <paragraph> (queue or fifo)

SEARCH

THE ACM DIGITAL LIBRARY

Advanced Search

Search Tips

Enter words, phrases or names below. Surround phrases or full names with double quotation marks.

Search within Results: **4,692** found

(bank <sentence> conflict*) <paragraph> (queue or fifo)

Clear result set

SEARCH

Desired Results:

must have **all** of the words or phrases

must have **any** of the words or phrases

must have **none** of the words or phrases

Name or Affiliation:

Authored by: ◉ all ○ any ○ none

Edited by: ◉ all ○ any ○ none

Reviewed by: ◉ all ○ any ○ none

Only search in:*

○ Title ○ Abstract ○ Review ◉ All Information

SEARCH

*Searches will be performed on all available information, including full text where available, unless specified above.

ISBN / ISSN: ◉ Exact ○ Expand

DOI: ◉ Exact ○ Expand

SEARCH

Published:

By: ◉ all ○ any ○ none

In: ◉ all ○ any ○ none

Since:

Month Year

Before:

February 2002

As: Any type of publication

Conference Proceeding:

Sponsored By:

Conference Location:

Conference Year:

yyyy

SEARCH

Classification: (CCS) ☐ **Primary Only**

Classified as: ◉ all ○ any ○ none

Subject Descriptor: ◉ all ○ any ○ none

Results must have accessible:

☐ Full Text ☐ Abstract ☐ Review

Published before February 2002

Found **33** of **33**

Sort results by: relevance
Display results: expanded form

Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 33

Result page: **1** 2 next

Relevance scale

1 A general framework for prefetch scheduling in linked data structures and its application to multi-chain prefetching
Seungryul Choi, Nicholas Kohout, Sumit Pamnani, Dongkeun Kim, Donald Yeung
May 2004 **ACM Transactions on Computer Systems (TOCS)**, Volume 22 Issue 2
Full text available: pdf(2.45 MB) Additional Information: full citation, abstract, references, index terms

Pointer-chasing applications tend to traverse composite data structures consisting of multiple independent pointer chains. While the traversal of any single pointer chain leads to the serialization of memory operations, the traversal of independent pointer chains provides a source of memory parallelism. This article investigates exploiting such *interchain memory parallelism* for the purpose of memory latency tolerance, using a technique called *multi--chain prefetching*. Previous work ...

Keywords: Data prefetching, memory parallelism, pointer-chasing code

2 Efficient use of memory bandwidth to improve network processor throughput
Jahangir Hasan, Satish Chandra, T. N. Vijaykumar
May 2003 **ACM SIGARCH Computer Architecture News , Proceedings of the 30th annual international symposium on Computer architecture**, Volume 31 Issue 2
Full text available: pdf(184.83 KB) Additional Information: full citation, abstract, references

We consider the efficiency of packet buffers used in packet switches built using network processors (NPs). Packet buffers are typically implemented using DRAM, which provides plentiful buffering at a reasonable cost. The problem we address is that a typical NP workload may be unable to utilize the peak DRAM bandwidth. Since the bandwidth of the packet buffer is often the bottleneck in the performance of a shared-memory packet switch, inefficient use of available DRAM bandwidth further reduces th ...

3 Banked multiported register files for high-frequency superscalar microprocessors
Jessica H. Tseng, Krste Asanović
May 2003 **ACM SIGARCH Computer Architecture News , Proceedings of the 30th annual international symposium on Computer architecture**, Volume 31 Issue 2
Full text available: pdf(142.29 KB) Additional Information: full citation, abstract, references, citings

Multiported register files are a critical component of high-performance superscalar microprocessors. Conventional multiported structures can consume significant power and die area. We examine the designs of banked multiported register files that employ multiple interleaved banks of fewer ported register cells to reduce power and area. Banked register files designs have been shown to provide sufficient bandwidth for a superscalar machine, but previous designs had complex control structures that w ...

4 The convex C240 architecture
M. Chastain, G. Gostin, J. Mankovich, S. Wallach
November 1988 **Proceedings of the 1988 ACM/IEEE conference on Supercomputing**

Full text available: pdf(961.68 KB) Additional Information: full citation, abstract, references, citings, index terms

The C240, a tightly coupled, shared memory, parallel/multi-processor, supports up to 4, 40-nanosecond ECL/CMOS Cray-like processors. Managed by a fully semaphored UNIX operating the C240, it can support up to 4 gigabytes of directly addressable physical memory. CONVEX proprietary compiler technology provides automatic vectorization and parallelization for FORTRAN, C, and ADA. Allocation of parallel threads to physical processors is managed by an innovation approach ASAP (Automatic Self-Allo ...

5 Register file and memory system design: Reducing register ports for higher speed and lower energy
Il Park, Michael D. Powell, T. N. Vijaykumar
November 2002 **Proceedings of the 35th annual ACM/IEEE international symposium on Microarchitecture**
Full text available: pdf(1.28 MB) Publisher Site Additional Information: full citation, abstract, references, citings, index terms

The key issues for register file design in high-performance processors are access time and energy. While previous work has focused on reducing the number of registers, we propose to reduce the number of register ports through two proposals, one for reads and the other for writes. For reads, we propose bypass hint to reduce register port requirements by avoiding unnecessary register file reads for cases where values are bypassed. Current processors are unable to avoid these unnecessary reads due ...

6 Multithreading and value prediction: Handling long-latency loads in a simultaneous multithreading processor
Dean M. Tullsen, Jeffery A. Brown
December 2001 **Proceedings of the 34th annual ACM/IEEE international symposium on Microarchitecture**
Full text available: pdf(1.17 MB) Publisher Site Additional Information: full citation, abstract, references, citings

Simultaneous multithreading architectures have been defined previously with fully shared execution resources. When one thread in such an architecture experiences a very long-latency operation, such as a load miss, the thread will eventually stall, potentially holding resources which other threads could be using to make forward progress.This paper shows that in many cases it is better to free the resources associated with a stalled thread rather than keep that thread ready to immediately begin ex ...

7 An instruction set and microarchitecture for instruction level distributed processing
Ho-Seop Kim, James E. Smith
May 2002 **ACM SIGARCH Computer Architecture News**, Volume 30 Issue 2
Full text available: pdf(1.08 MB) Publisher Site Additional Information: full citation, abstract, references, citings, index terms

An instruction set architecture (ISA) suitable for future microprocessor design constraints is proposed. The ISA has hierarchical register files with a small number of accumulators at the top. The instruction stream is divided into chains of dependent instructions (strands) where intra-strand dependences are passed through the accumulator. The general-purpose register file is used for communication between strands and for holding global values that have many consumers.A microarchitecture to supp ...

8 A dynamic multithreading processor
Haitham Akkary, Michael A. Driscoll
November 1998 **Proceedings of the 31st annual ACM/IEEE international symposium on Microarchitecture**
Full text available: pdf(2.67 MB) Additional Information: full citation, references, citings, index terms

9 The specification of a new Manchester Dataflow machine
uhir In g mi hn ey

June 1986 **Proceedings of the 3rd international conference on Supercomputing**

Full text available: pdf(996.06 KB) Additional Information: full citation, abstract, references, citings, index terms

The Present Manchester Dataflow Machine is constructed using standard TTL technology and is not designed with raw processing power in mind. This paper discusses the issues raised in the redesign of the machine using supercomputer technology. The resulting machine structure is presented in some detail, together with the initial results of simulation which indicate that a performance of hundreds of megaflops appears readily achievable.

10 A hardware accelerator for maze routing
Y. Won, S. Sahni, Y. El-ziq
October 1987 **Proceedings of the 24th ACM/IEEE conference on Design automation**

Full text available: pdf(871.73 KB) Additional Information: full citation, abstract, references, index terms

A hardware accelerator for the maze routing problem is developed. This accelerator consists of three 3 stage pipelines. Banked memory is used to avoid memory read/write conflicts and obtain maximum efficiency.

11 Processor-memory coexploration using an architecture description language
Prabhat Mishra, Mahesh Mamidipaka, Nikil Dutt
February 2004 **ACM Transactions on Embedded Computing Systems (TECS)**, Volume 3 Issue 1

Full text available: pdf(201.88 KB) Additional Information: full citation, abstract, references, index terms

Memory represents a major bottleneck in modern embedded systems in terms of cost, power, and performance. Traditionally, memory organizations for programmable embedded systems assume a fixed cache hierarchy. With the widening processor--memory gap, more aggressive memory technologies and organizations have appeared, allowing customization of a heterogeneous memory architecture tuned for specific target applications. However, such a processor--memory coexploration approach critically needs the ab ...

Keywords: Processor--memory codesign, architecture description language, design space exploration, memory exploration

12 Design and Implementation of High-Performance Memory Systems for Future Packet Buffers
Jorge García, Jesús Corbal, Llorenç Cerdà, Mateo Valero
December 2003 **Proceedings of the 36th Annual IEEE/ACM International Symposium on Microarchitecture**

Full text available: pdf(348.55 KB) Publisher Site Additional Information: full citation, abstract

In this paper we address the design of a future high-speedrouter that supports line rates as high as OC-3072 (160 Gb/s),around one hundred ports and several service classes. Buildingsuch a high-speed router would raise many technological problems,one of them being the packet buffer design, mainly becausein router design it is important to provide worst-case bandwidthguarantees and not just average-case optimizations.A previous packet buffer design provides worst-case bandwidthguarantees by using ...

13 The effects of input/output activity on the average instruction execution time of a real-time computer system
Salvatore C. Catania
December 1969 **Proceedings of the third conference on Applications of simulation**

Full text available: pdf(591.64 KB) Additional Information: full citation, abstract, index terms

The paper describes the results of a GPSS simulation of a real-time computer system undertaken to determine the effects of specific types of Input/Output (I/O) activity on the average instruction execution time (AIET) of the system's central processor. The types of I/O activities studied are communications I/0 and references to mass storage devices. Various memory bank configurations and different memory bank referencing schemes are considered in an attempt to minimize the effects of the I ...

14 Embedded hardware design case studies: A fully-programmable memory management system optimizing queue handling at multi gigabit rates
G. Kornaros, I. Papaefstathiou, A. Nikologiannis, N. Zervos
June 2003 **Proceedings of the 40th conference on Design automation**

Full text available: pdf(236.13 KB) Additional Information: full citation, abstract, references, index terms

Two of the main bottlenecks when designing a network embedded system are very often the memory bandwidth and its capacity. This is mainly due to the extremely high speed of the state-of-the-art network links and to the fact that in order to support advanced quality of service (QoS), per-flow queueing is desirable. In this paper we describe the architecture of a memory manager that can provide up to 10Gbs of aggregate throughput while handling 512K queues. The presented system supports a complete ...

Keywords: memory management, network processor

15 An aperiodic storage scheme to reduce memory conflicts in vector processors
S. Weiss
April 1989 **ACM SIGARCH Computer Architecture News , Proceedings of the 16th annual international symposium on Computer architecture**, Volume 17 Issue 3

Full text available: pdf(649.82 KB) Additional Information: full citation, abstract, references, citings, index terms

One of the most noticeable differences between the CRAY-2 and its predecessors, the CRAY-1 and the CRAY X-MP, is a significantly longer memory path. This is a consequence of increasing the size of the memory at the expense of the bank access time. With a longer memory path, the impact of bank conflicts becomes more apparent. In this paper we study a storage strategy for vector processors that has the following properties: (1) it is aperiodic, (2) it tends to distribute references more unifo ...

16 A processor architecture for horizon
M. R. Thistle, B. J. Smith
November 1988 **Proceedings of the 1988 ACM/IEEE conference on Supercomputing**

Full text available: pdf(970.44 KB) Additional Information: full citation, abstract, references, citings, index terms

Horizon is a scalable shared-memory Multiple Instruction stream - Multiple Data stream (MIMD) computer architecture independently under study at the Supercomputing Research Center (SRC) and Tera Computer Company. It is composed of a few hundred identical scalar processors and a comparable number of memories, sparsely embedded in a three-dimensional nearest-neighbor network. Each processor has a horizontal instruction set that can issue up to three floating point operations per cycle without ...

17 Session 9: hardware performance: Cache performance in vector supercomputers
L. I. Kontothanassis, R. A. Sugumar, G. J. Faanes, J. E. Smith, M. L. Scott
November 1994 **Proceedings of the 1994 ACM/IEEE conference on Supercomputing**

Full text available: pdf(780.13 KB) Additional Information: full citation, abstract, references

Traditional supercomputers use a flat multi-bank SRAM memory organization to supply high bandwidth at low latency. Most other computers use a hierarchical organization with a small SRAM cache and slower, cheaper DRAM for main memory. Such systems rely heavily on data locality for achieving optimum performance. This paper evaluates cache-based memory systems for vector supercomputers. We develop a simulation model for a cache-based version of the Cray Research C90 and use the NAS parallel benchma ...

18 Session 7: Tradeoffs in power-efficient issue queue design
Alper Buyuktosunoglu, David H. Albonesi, Pradip Bose, Peter W. Cook, Stanley E. Schuster
August 2002 **Proceedings of the 2002 international symposium on Low power electronics and design**

Full text available: pdf(92.97 KB) Additional Information: full citation, abstract, references, index terms

A major consumer of microprocessor power is the issue queue. Several microprocessors, including the Alpha 21264 and POWER4™, use a compacting latch-based issue queue design which h s the dv nt ge f simplicity f design nd verific ti n. he dis dv nt ge f this

structure, however, is its high power dissipation.In this paper, we explore different issue queue power optimization techniques that vary not only in their performance and power characteristics, but in how much they deviate ...

Keywords: adaptation, banking, compacting, issue queue, low-power, microarchitecture, non-compacting

19 Synchronizing processors through memory requests in a tightly coupled multiprocessor

A. Seznec, Y. Jégou

May 1988 **ACM SIGARCH Computer Architecture News , Proceedings of the 15th Annual International Symposium on Computer architecture**, Volume 16 Issue 2

Full text available: pdf(927.94 KB) Additional Information: full citation, abstract, references, citings, index terms

To satisfy the growing need for computing power, a high degree of parallelism will be necessary in future supercomputers. Up to the late 70s, supercomputers were either multiprocessors (SIMD-MIMD) or pipelined monoprocessors. Current commercial products combine these two levels of parallelism. Effective performance will depend on the spectrum of algorithms which is actually run in parallel. In a previous paper [Je86], we have presented the DSPA processor, a pipeline processor whi ...

20 Efficient synchronization of multiprocessors with shared memory

Clyde P. Kruskal, Larry Rudolph, Marc Snir

October 1988 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 10 Issue 4

Full text available: pdf(1.78 MB) Additional Information: full citation, abstract, references, citings, index terms, review

A new formalism is given for read-modify-write (RMW) synchronization operations. This formalism is used to extend the memory reference combining mechanism introduced in the NYU Ultracomputer, to arbitrary RMW operations. A formal correctness proof of this combining mechanism is given. General requirements for the practicality of combining are discussed. Combining is shown to be practical for many useful memory access operations. This includes memory updates of the form mem_ ...